■ ECON 252 Financial Markets - 02 The Universal Principle of Risk Management_ Pooling and the Hedging of Risks

URL: http://oyc.yale.edu/economics/financial-markets/content/sessions/session-2-the-universal-principle-of-risk

Lecture Briefing

ECON 252: Financial Markets

Lecture 2 - The Universal Principle of Risk Management: Pooling and the Hedging of Risks	<< previous session | next session >>

Overview:

Statistics and mathematics underlie the theories of finance. Probability Theory and various distribution types are important to understanding finance. Risk management, for instance, depends on tools such as variance, standard deviation, correlation, and regression analysis. Financial analysis methods such as present values and valuing streams of payments are fundamental to understanding the time value of money and have been in practice for centuries.

Reading assignment:

Jeremy Siegel, *Stocks for the Long Run*, chapter 1 and Appendix 2 (p. 12)

Class lecture:

- **視訊檔案下載 (亦可線上觀看)**:
  ■ ECON 252 Financial Markets - 2 The Universal Principle of Risk Management: Pooling and the Hedging of Risks
- **我儲存視訊檔案的網頁**:
  (http://www.archive.org/details/Econ2522008fall02-TheUniversa)
- **ALPHA VIDEO ■ Lex Economie 部落格處本講次之專欄網址**:
  ■ ECON 252 Financial Markets - 2 The Universal Principle of Risk Management: Pooling and the Hedging of Risks

Financial Markets: Lecture 2 Transcript

URL:
http://oyc.yale.edu/economics/financial-markets/content/transcripts/transcript-2-the-universal-principle-of-risk

January 16, 2008	<< back

Professor Robert Shiller: Today I want to spend--The title of today's lecture is: The Universal Principle of Risk Management, Pooling and the Hedging of Risk. What I'm really referring to is what I think is the very original, the deep concept that underlies theoretical finance--I wanted to get that first. It really is probability theory and the idea of spreading risk through risk pooling. So, this idea is an intellectual construct that appeared at a certain point in history and it has had an amazing number of applications and finance is one of these. Some of you--This incidentally will be a more technical of my lectures and it's a little bit unfortunate that it comes early in the semester. For those of you who have had a course in probability and statistics, there will be nothing new here. Well, nothing in terms of the math. The probability theory is new. Others though, I want to tell you that it doesn't--if you're shopping--I had a student come by yesterday and ask--he's a little rusty in his math skills--if he should take this course. I said, "Well if you can understand tomorrow's lecture--that's today's lecture--then you should have no problem."

I want to start with the concept of probability. Do you know what a probability is? We attach a probability to an event. What is the probability that the stock market will go up this year? I would say--my personal probability is .45. That's because I'm a bear but-- Do you know what that means? That 45 times out of 100 the stock market will go up and the other 55 times out of 100 it will stay the same or go down. That's a probability. Now, you're familiar with that concept, right? If someone says the probability is .55 or .45, well you know what that means. I want to emphasize that it hasn't always been that way and that probability is really a concept that arose in the 1600s. Before that, nobody ever said that.

Ian Hacking, who wrote a history of probability theory, searched through world literature for any reference to a probability and could find none anywhere before 1600. There was an intellectual leap that occurred in the seventeenth century and it became very fashionable to talk in terms of probabilities. It spread throughout the world--the idea of quoting probabilities. But it was--It's funny that such a simple idea hadn't been used before. Hacking points out that the word probability--or probable--was already in the English language. In fact, Shakespeare used it, but what do you think it meant? He gives an example of a young woman, who was describing a man that she liked, and she said, I like him very much, I find him very probable. What do you think she means? Can someone answer that? Does anyone know Elizabethan English well enough to tell me? What is a probable young man? I'm asking for an answer. It sounds like people have no idea. Can anyone venture a guess? No one wants to venture a guess?

Student: fertile?

Professor Robert Shiller: That he can father children? I don't think that's what she meant but maybe. No, what apparently she meant is trustworthy. That's a very important quality in a person I suppose.

So, if something is probable you mean that you can trust it and so probability means trustworthiness. You can see how they moved from that definition of probability to the current definition. But Ian Hacking, being a good historian, thought that someone must have had some concept of probability going before, even if they didn't quote it as a number the way--it must have been in their head or in their idea. He searched through world literature to try to find some use of the term that preceded the 1600s and he concluded that there were probably a number of people who had the idea, but they didn't publish it, and it never became part of the established literature partly because, he said, throughout human history, there has been a love of gambling and probability theory is extremely useful if you are a gambler. Hacking believes that there were many gambling theorists who invented probability theory at various times in history but never wrote it down and kept it as a secret.

He gives an example--I like to--he gives an example from a book that--or it's a collection--I think, a collection of epic poems written in Sanskrit that goes back--it was actually written over a course of 1,000 years and it was completed in the fourth century. Well, there's a story--there's a long story in the Mahabarahta about an emperor called Nala and he had a wife named Damayanti and he was a very pure and very good person. There was an evil demon called Kali who hated Nala and wanted to bring his downfall, so he had to find a weakness of Nala. He found finally some, even though Nala was so pure and so perfect--he found one weakness and that was gambling. Nala couldn't resist the opportunity to gamble; so the evil demon seduced him into gambling aggressively. You know sometimes when you're losing and you redouble and you keep hoping to win back what you've lost? In a fit of gambling, Nala finally gambled his entire kingdom and lost--it's a terrible story--and Nala then had to leave the kingdom and his wife. They wandered for years. He separated from her because of dire necessity.

They were wandering in the forests and he was in despair, having lost everything. But then he meets someone by the name of--we have Nala and he meets this man, Rituparna, and this is where a probability theory apparently comes in. Rituparna tells Nala that he knows the science of gambling and he will teach it to Nala, but that it has to be done by whispering it in his ear because it's a deep and extreme secret. Nala is skeptical. How does Rituparna know how to gamble? So Rituparna tries to prove to him his abilities and he says, see that tree there, I can estimate how many leaves there are on that tree by counting leaves on one branch. Rituparna looked at one branch and estimated the number of leaves on the tree, but Nala was skeptical. He stayed up all night and counted every leaf on the tree and it came very close to what Rituparna said; so he--the next morning--believed Rituparna. Now this is interesting, Hacking says, because it shows that sampling theory was part of Nala's theory. You don't have to count all the leaves on the tree, you can take a sample and you count that and then you multiply.

Anyway, the story ends and Nala goes back and is now armed with probability theory, we assume. He goes back and gambles again, but he has nothing left to wager except his wife; so he puts her and gambles her. But remember, now he knows what he's doing and so he really wasn't gambling his wife--he was really a very pure and honorable man. So he won back the entire kingdom and that's the ending.

Anyway, that shows that I think probability theory does have a long history, but--it not being an intellectual discipline--it didn't really inform a generation of finance theory. When you don't have a theory, then you don't have a way to be rigorous. So, it was in

the 1600s that probability theory started to get written down as a theory and many things then happened in that century that, I think, are precursors both to finance and insurance. One was in the 1600s when people started constructing life tables. What is a life table? It's a table showing the probability of dying at each age, for each age and sex. That's what you need to know if you're going to do life insurance. So, they started to do collecting of data on mortality and they developed something called actuarial science, which is estimating the probability of people living. That then became the basis for insurance. Actually, insurance goes back to ancient Rome in some form. In ancient Rome they had something called burial insurance. You could buy a policy that protected you against your family not having the money to bury you if you died. In ancient culture people worried a great deal about being properly buried, so that's an interesting concept. They were selling that in ancient Rome; but you might think, but why just for burial? Why don't you make it into full-blown life insurance? You kind of wonder why they didn't. I think maybe it's because they didn't have the concepts down. In Renaissance Italy they started writing insurance policies--I read one of the insurance policies, it's in the Journal of Risk and Insurance--and they translate a Renaissance insurance policy and it's very hard to understand what this policy was saying. I guess they didn't have our language, they didn't--they were intuitively halfway there but they couldn't express it, so I think the industry didn't get really started. I think it was the invention of probability theory that really started it and that's why I think theory is very important in finance.

Some people date fire insurance with the fire of London in 1666. The whole city burned down, practically, in a terrible fire and fire insurance started to proliferate right after that in London. But you know, you kind of wonder if that's a good example for fire insurance because if the whole city burns down, then insurance companies would go bankrupt anyway, right? London insurance companies would because the whole concept of insurance is pooling of independent probabilities. Nonetheless, that was the beginning. We're also going to recognize, however, that insurance got a slow start because--I believe it is because--people could not understand the concept of probability. They didn't have the concept firmly in mind. There are lots of aspects to it. In order to understand probability, you have to take things as coming from a random event and people don't clearly have that in their mind from an intuitive standpoint. They have maybe a sense that I can influence events by willing or wishing and if I think that--if I have kind of a mystical side to me, then probabilities don't have a clear meaning. It has been shown that even today people seem to think that. They don't really take, at an intuitive level, probabilities as objective. For example, if you ask people how much they would be willing to bet on a coin toss, they will typically bet more if they can toss the coin or they will bet more if the coin hasn't been tossed yet. It could have been already tossed and concealed. Why would that be? It might be that there's just some intuitive sense that I can--I don't know--I have some magical forces in me and I can change things.

The idea of probability theory is that no, you can't change things, there are all these objective laws of probability out there that guide everything. Most languages around the world have a different word for luck and risk--or luck and fortune. Luck seems to mean something about you: like I'm a lucky person. I don't know what that means--like God or the gods favor me and so I'm lucky or this is my lucky day. Probability theory is really a movement away from that. We then have a mathematically rigorous discipline.

Now, I'm going to go through some of the terms of probability and--this will be review for

many of you, but it will be something that we're going to use in the--So I'll use the symbol *P* or I can sometimes write it out as *prob* to represent a probability. It is always a number that lies between zero and one, or between 0% and 100%. "Percent" means divided by 100 in Latin, so 100% is one. If the probability is zero that means the event can't happen. If the probability is one, it means that it's certain to happen. If the probability is--Can everyone see this from over there? I can probably move this or can't I? Yes, I can. Now, can you now--you're the most disadvantaged person and you can see it, right? So that's the basic idea.

One of the first principles of probability is the idea of independence. The idea is that probability measures the likelihood of some outcome. Let's say the outcome of an experiment, like tossing a coin. You might say the probability that you toss a coin and it comes up heads is a half, because it's equally likely to be heads and tails. Independent experiments are experiments that occur without relation to each other. If you toss a coin twice and the first experiment doesn't influence the second, we say they're independent and there's no relation between the two.

One of the first principles of probability theory is called the multiplication rule. That says that if you have independent probabilities, then the probability of two events is equal to the product of their probabilities. So, the _Prob(A_ and _B) = Prob(A)*Prob(B)._ That wouldn't hold if they're not independent. The theory of insurance is that ideally an insurance company wants to insure independent events. Ideally, life insurance is insuring people--or fire insurance is insuring people--against independent events; so it's not the fire of London. It's the problem that sometimes people knock over an oil lamp in their home and they burn their own house down. It's not going to burn any other houses down since it's just completely independent of anything else. So, the probability that the whole city burns down is infinitesimally small, right? This will generalize to probability of *A* and *B* and *C* equals the probability of *A* times the probability of *B* times the probability of *C* and so on. If the probability is 1 in 1,000 that a house burns down and there are 1,000 houses, then the probability that they all burn down is 1/1000 to the 1000th power, which is virtually zero. So insurance companies then–Basically, if they write a lot of policies, then they have virtually no risk. That is the fundamental idea that may seem simple and obvious to you, but it certainly wasn't back when the idea first came up.

Incidentally, we have a problem set, which I want you to start today and it will be due not in a week this time, because we have Martin Luther King Day coming up, but it will be due the Monday following that.

If you follow through from the independent theory, there's one of the basic relations in probability theory--it's called the binomial distribution. I'm not going to spend a whole lot of time on this but it gives the probability of *x* successes in *n* trials or, in the case of insurance *x,* if you're insuring against an accident, then the probability that you'll get *x* accidents and *n* trials. The binomial distribution gives the probability as a function of *x* and it's given by the formula where *P* is the probability of the accident: PX (1-P)N-X [n! /(n-x)!]. That is the formula that insurance companies use when they have independent probabilities, to estimate the likelihood of having a certain number of accidents. They're concerned with having too many accidents, which might exhaust their reserves. An insurance company has reserves and it has enough reserves to cover them for a certain number of accidents. It uses the binomial distribution to calculate the probability of getting any specific number of accidents. So, that is the binomial distribution. I'm not

going to expand on this because I can't get into--This is not a course in probability theory but I'm hopeful that you can see the formula and you can apply it. Any questions? Is this clear enough? Can you read my handwriting?

Another important concept in probability theory that we will use a lot is expected value, the mean, or average--those are all roughly interchangeable concepts. We have expected value, mean or average. We can define it in a couple of different ways depending on whether we're talking about sample mean or population mean. The basic definition--the expected value of some random variable *x--E(x)*--I guess I should have said that a random variable is a quantity that takes on value. If you have an experiment and the outcome of the experiment is a number, then a random variable is the number that comes from the experiment. For example, the experiment could be tossing a coin; I will call the outcome *heads* the number one, and I'll call the outcome *tails* the number zero, so I've just defined a random variable. You have discrete random variables, like the one I just defined, or there are also--which take on only a finite number of values--and we have continuous random variables that can take on any number of values along a continuum. Another experiment would be to mix two chemicals together and put a thermometer in and measure the temperature. That's another invention of the 1600s, by the way--the thermometer. And they learned that concept--perfectly natural to us--temperature. But it was a new idea in the 1600s. So anyway, that's continuous, right? When you mix two chemicals together, it could be any number, there's an infinite number of possible numbers and that would be continuous.

For discrete random variables, we can define the expected value, or µx --that's the Greek letter *mu*--as the summation i = 1 to infinity of. [P(x=xi) times (xi)]. I have it down that there might be an infinite number of possible values for the random variable *x*. In the case of the coin toss, there are only two, but I'm saying in general there could be an infinite number. But they're accountable and we can list all possible values when they're discrete and form a probability weighted average of the outcomes. That's called the expected value. People also call that the mean or the average. But, note that this is based on theory. These are probabilities. In order to compute using this formula you have to know the true probabilities. There's another formula that applies for a continuous random variables and it's the same idea except that--I'll also call it µx, except that it's an integral. We have the integral from minus infinity to plus infinity of *F(x)*x*dx*, and that's really--you see it's the same thing because an integral is analogous to a summation. Those are the two population definitions. *F(x)* is the continuous probability distribution for *x*. That's different when you have continuous values--you don't have *P (x = xi)* because it's always zero. The probability that the temperature is exactly 100° is zero because it could be 100.0001° or something else and there's an infinite number of possibilities. We have instead what's called a probability density when we have continuous random variables. You're not going to need to know a lot about this for this course, but this is--I wanted to get the basic ideas down. These are called population measures because they refer to the whole population of possible outcomes and they measure the probabilities. It's the truth, but there are also sample means. When you get--this is Rituparna, counting the leaves on a tree--you can estimate, from a sample, the population expected values. The population mean is often written "*x-bar*." If you have a sample with *n* observations, it's the summation *i* = 1 to *n* of *xi/n*--that's the average. You know that formula, right? You count *n* leaves--you count the number of leaves. You have

n branches on the tree and you count the number of leaves and sum them up. One would be--I'm having a little trouble putting this into the Rituparna story, but you see the idea. You know the average, I assume. That's the most elementary concept and you could use it to estimate either a discreet or continuous expected value.

In finance, there's often reference to another kind of average, which I want to refer you to and which, in the Jeremy Siegel book, a lot is made of this. The other kind of average is called the geometric average. We'll call that--I'll only show the sample version of it G(x) = the product *i* = 1 to *n* of (xi)^(1/n). Does everyone--Can you see that? Instead of summing them and dividing by M, I multiply them all together and take the nth root of them. This is called the geometric average and it's used only for positive numbers. So, if you have any negative numbers you'd have a problem, right? If you had one negative number in it, then the product would be a negative number and, if you took a root of that, then you might get an imaginary number. We don't want to use it in that case.

There's an appendix to one of the chapters in Jeremy Siegel's book where he says that one of the most important applications of this theory is to measure how successful an investor is. Suppose someone is managing money. Have they done well? If so, you would say, "Well, they've been investing money over a number of different years. Let's take the average over all the different years." Suppose someone has been investing money for *n* years and *xi* is the return on the investment in a given year. What is their average performance? The natural thing to do would be to average them up, right? But Jeremy says that maybe that's not a very good thing to do. What he says you should do instead is to take the geometric average of gross returns. The return on an investment is how much you made from the investment as a percent of the money invested. The gross return is the return plus one. The worst you can ever do investing is lose all of your investment--lose 100%. If we add one to the return, then you've got a number that's never negative and we can then use geometric returns.

Jeremy Siegel says that in finance we should be using geometric and not arithmetic averages. Why is that? Well I'll tell you in very simple terms, I think. Suppose someone is investing your money and he announces, I have had very good returns. I have invested and I've produced 20% a year for nine out of the last ten years. You think that's great, but what about the last year. The guy says, "Oh I lost 100% in that year." You might say, "Alright, that's good." I would add up 20% a year for nine years and than put in a zero–no, 120 because it's gross return for nine years--and put in a zero for one year. Maybe that doesn't look bad, right? But think about it, if you were investing your money with someone like that, what did you end up with? You ended up with nothing. If they have one year when they lose everything, it doesn't matter how much they made in the other years. Jeremy says in the text that the geometric return is always lower than the arithmetic return unless all the numbers are the same. It's a less optimistic version. So, we should use that, but people in finance resist using that because it's a lower number and when you're advertising your return you want to make it look as big as possible.

We also need some measure of--We've been talking here about measures of central tendency only and in finance we need, as well, measures of dispersion, which is how much something varies. Central tendency is a measure of the center of a probability distribution of the--Central tendency is a measure--Variance is a measure of how much things change from one observation to another. We have variance and it's often represented by σ^2, that's the Greek letter sigma, lower case, squared. Or, especially

when talking about estimates of the variance, we sometimes say S^2 or we say *standard deviation*^{2}. The standard deviation is the square root of the variance. For population variance, the variance of some random variable *x* is defined as the summation i = 1 to infinity of the *Prob (x = xi)* times *(xi - μx)2*. So mu is the mean--we just defined it of *x*--that's the expectation of *x* or also *E(x)*, so it's the probability weighted average of the squared deviations from the mean. If it moves a lot--either way from the mean--then this number squared is a big number. The more *x* moves, the bigger the variance is.

There's also another variance measure, which we use in the sample--or also Var is used sometimes--and this is $\sum^2$. There's also another variance measure, which is for the sample. When we have *n* observations it's just the summation i = 1 to *n* of *(x - x bar)²/n*. That is the sample variance. Some people will divide by *n–1*. I suppose I would accept either answer. I'm just keeping it simple here. They divide by *n-1* to make it an unbiased estimator of the population variance; but I'm just going to show it in a simple way here. So you see what it is--it's a measure of how much *x* deviates from the mean; but it's squared. It weights big deviations a lot because the square of a big number is really big. So, that's the variance.

So, that completes central tendency and dispersion. We're going to be talking about these in finance in regards to returns because--generally the idea here is that we want high returns. We want a high expected value of returns, but we don't like variance. Expected value is good and variance is bad because that's risk; that's uncertainty. That's what this whole theory is about: how to get a lot of expected return without getting a lot of risk.

Another concept that's very basic here is covariance. Covariance is a measure of how much two variables move together. Covariance is--we'll call it--now we have two random variables, so I'll just talk about it in a sample term. It's the summation *i* = 1 to *n of [(x – x-bar)* times *(y – y-bar)]/n*. So *x* is the deviation for the *i*-subscript, meaning we have a separate *xi* and yi for each observation. So we're talking about an experiment when you generate--Each experiment generates both an *x* and a *y* observation and we know when *x* is high, *y* also tends to be high, or whether it's the other way around. If they tend to move together, when *x* is high and *y* is high together at the same time, then the covariance will tend to be a positive number. If when *x* is low, *y* also tends to be low, then this will be negative number and so will this, so their product is positive. A positive covariance means that the two move together. A negative covariance means that they tend to move opposite each other. If *x* is high relative to *x-bar*--this is positive--then *y* tends to be low relative to its mean *y-bar* and this is negative. So the product would be negative. If you get a lot of negative products, that makes the covariance negative.

Then I want to move to correlation. So this is a measure--it's a scaled covariance. We tend to use the Greek letter *rho*. If you were to use Excel, it would be *correl* or sometimes I say *corr*. That's the correlation. This number always lies between -1 and +1. It is defined as rho= [cov(xiyi)/SxSy] That's the correlation coefficient. That has kind of almost entered the English language in the sense that you'll see it quoted occasionally in newspapers. I don't know how much you're used to it--Where would you see that? They would say there is a low correlation between SAT scores and grade point averages in college, or maybe it's a high correlation. Does anyone know what it is? But you could estimate the corr--it's probably positive. I bet it's way below one, but it has some correlation, so maybe it's .3. That would mean that people who have high SAT

scores tend to get higher grades. If it were negative--it's very unlikely that it's negative--it couldn't be negative. It couldn't be that people who have high SAT scores tend to do poorly in college. If you quantify how much they relate, then you could look at the correlation.

I want to move to regression. This is another concept that is very basic to statistics, but it has particular use in finance, so I'll give you a financial example. The concept of regression goes back to the mathematician Gauss, who talked about fitting a line through a scatter of points. Let's draw a line through a scatter of points here. I want to put down on this axis the return on the stock market and on this axis I want to put the return on one company, let's say Microsoft. I'm going to have each observation as a year. I shouldn't put down a name of a company because I can't reproduce this diagram for Microsoft. Let's not say Microsoft, let's say Shiller, Inc. There's no such company, so I can be completely hypothetical. Let's put zero here because these are not gross returns these are returns, so they're often negative. Suppose that in a given year--and say this is minus five and this is plus five, this is minus five and this is plus five--Suppose that in the first year in our sample, the company Shiller, Inc. and the market both did 5%. That puts a point right there at five and five. In another year, however, the stock market lost 5% and Shiller, Inc. lost 7%. We would have a point, say, down here at five and seven. This could be 1979, this could be 1980, and we keep adding points so we have a whole scatter of points. It's probably upward sloping, right? Probably when the overall stock market does well so does Shiller, Inc.

What Gauss did was said, let's fit a line through the point--the scatter of points--and that's called the regression line. He chose the line so that--this is Gauss--he chose the line to minimize the sum of squared distances of the points from the lines. So these distances are the lengths of these line segments. To get the best fitting line, you find the line that minimizes the sum of squared distances. That's called the regression line and the intercept is called *alpha*--there's *alpha*. And the slope is called *beta*. That may be a familiar enough concept to you, but in the context of finance this is a major concept. The way I've written it, the *beta* of Shiller, Inc. is the slope of this line. The *alpha* is just the intercept of this curve. We can also do this with excess returns. I will get to this later, where I have the return minus the interest rate on this axis and the market return minus the interest rate on this axis. In that case, *alpha* is a measure of how much Shiller, Inc. outperforms. We'll come back to this, but *beta* of the stock is a measure of how much it moves with the market and the *alpha* of a stock is how much it outperforms the market. We'll have to come back to that--these are basic concepts.

I want to--another concept--I guess I've just been implicit in what I have--There's a distribution called the normal distribution and that is--I'm sure you've heard of this, right? If you have a distribution that looks like this--it's bell-shaped--this is *x* and--I have to make it look symmetric which I may not be able to do that well--this is *f(x),* the normal distribution. $f(x) = [1/(\sqrt{}\ (2\pi)\sigma)]$ times e to minus $[(x-\mu)2\ /\ 2\sigma]$. It's a famous formula, which is due to Gauss again. We often assume in finance that random variables, such as returns, are normally distributed. This is called the normal distribution or the Gaussian distribution--it's a continuous distribution. I think you've heard of this, right? This is high school raw material. But I want to emphasize that there are also other bell-shaped curves. This is the most famous bell-shaped curve, but there are other ones with different mathematics.

A particular interest in finance is fat-tailed alternatives. It could be that a random distribution--I don't have colored chalk here I don't think, so I will use a dash line to represent the fat-tailed distribution. Suppose the distribution looks like this. Then I have to try to do that on the other side, as symmetrically as I can. These are the tails of the distribution; this is the right tail and this is the left tail. You can see that the dash distribution I drew has more out in the tails, so we call it fat-tailed. This refers to random variables that have fat-tailed distributions--random variables that occasionally give you really big outcomes. You have a chance of being way out here with a fat-tailed distribution. It's a very important observation in finance that returns on a lot of speculative assets have fat-tailed distributions. That means that you can go through twenty years of a career on Wall Street and all you've observed is observations in the central region. So you feel that you know pretty well how things behave; but then, all of a sudden, there's something way out here. This would be good luck if you were long and now suddenly you got a huge return that you would not have thought was possible since you've never seen it before. But you can also have an incredibly bad return. This complicates finance because it means that you never know. You never have enough experience to get through all these things. It's a big complication in finance.
My friend Nassim Talib has just written a book about it called--maybe I'll talk about that--called *The Black Swan*. It's about how so many plans in finance are messed up by rare events that suddenly appear out of nowhere. He called it *The Black Swan* because if you look at swans, they're always white. You've never seen a black swan. So, you end up going through life assuming that there are no black swans. But, in fact, there are and you might finally see one. You don't want to predicate making complicated gambles under the assumption that they don't exist. Talib, who's a Wall Street professional, talks about these black swans as being the real story of finance.
Now. I want to move away from statistics and talk about present values, which is another concept in finance that is fundamental. And so, let me--And then this will conclude today's lecture. What is a present value? This isn't really statistics anymore, but it's a concept that I want to include in this lecture. People in business often have claims on future money, not money today. For example, I may have someone who promises to pay me $1 in one year or in two years or three years. The present value is what that's worth today. I may have an "IOU" from someone or I may own a bond from someone that promises to pay me something in a year or two years. According to a time-honored tradition in finance, it says that it's a promise to pay $1, but it's not worth $1 today. It must be worth less than $1. What you could do hundreds of years ago--and can still do it today--was go to a bank and present this bond or IOU and say, "What will you give me for it?" The bank will discount it. Sometimes we say "present discounted value." The banker will say, "Well you have $1 a year from now, but that's a year from now, so I won't give you $1 now. I'll give you the present discounted value for it."
Now, I'm going to abstract from risk. Let's assume that we know that this thing is going to be paid, so it's a matter of simple time. Of course, the banker isn't going to give you $1 for something that is paying $1 in a year because the banker knows that $1 could be invested at the interest rate. Let's say the interest rate is r and that would be a number like .05, which is 5%, which is five divided by one hundred. Then the present value of $1--The present PDV or PV of $1 = $1/(1+$r$). That's because the banker is thinking, if I have this amount right now and I invest it for one year, then what do I have. I have (1 +

r)*(1/1+r). It's $1, so that works out exactly right. You have to discount something that's one period in the future by dividing it by 1+r. This is the present value of $1 in one time period, which I'm taking to be a year. It doesn't have to be a year. The interest rate has units of time, so I have to specify a time period over which I'm measuring an interest rate. Typically it's a year. If it's a one-year interest rate, the time period is one year, and the present value of $1 in one time period is given by this: the present value of $1 in n periods is 1/(1+r)n and that's all there is to this.

I want to talk about valuing streams of payments. Suppose someone has a contract that promises to pay an amount each period over a number of years. We have formulas for these present values and these formulas are well known. I'm just going to go through them rather quickly here. The simplest thing is something called a consol or perpetuity. A perpetuity is an asset or a contract that pays a fixed amount of money each time period, forever. We call them consols because, in the early 1700s, the British Government issued what they called consols or consolidated debt of the British Crown that paid a certain amount of pound sterling every six months forever. You may say, what audacity for the British Government to promise to pay anything forever. Will they be around forever? Well as far as you're concerned, it's as good as forever, right? Maybe someday the British--United Kingdom--something will happen to it, it will fall apart or change; but that is so distant in the future that we can disregard that, so we'll take that as forever. Anyway, the government might buy them back too, so who cares if it isn't forever. Let's just talk about it as forever.

Let's say this thing pays one pound a period forever. What is the present value of that? Well, the first--each payment we'll call a coupon--so it pays one pound one year from now. Let's say it's one year just to simplify things. It pays another pound two years from now, it pays another pound three years from now. The present value is equal to--remember it starts one year from now under assumption--we could do it differently but I'm assuming one year now. The present value is 1/(1+r) for the first year; plus for the second year it's $1/(1+r)^2$; for the third year it's $1/(1+r)^3$; and that goes on forever. That's an infinite series and you know how to sum that, I think. I'll tell you what it is: it's 1/r, or it would be £1/r . Generally, if it pays c dollars for every period, the present value is c/r. That's the formula for the present value of a consol. That's one of the most basic formulas in finance. The interesting thing is that it means that the value of consol moves inversely to the interest rate. The British Government issued those consols in the early 1700s and, while they were refinanced in the late nineteenth century, they're still there. If you want to go out and buy one, you can get on your laptop right after this lecture and buy one of them. Then you've got something that will pay you something forever. But you're going to know that the value of that in the market moves opposite with interest rates. So, if interest rates go down, the value goes up; if interest rates go up, the value of your investment goes down.

Another formula is--what if the consol doesn't pay--I'm sorry, the next thing is a growing consol. I'm calling it a growing consol even though the British consols didn't grow. Let's say that the British Government didn't say that they'll pay one pound per year, but it'll be one pound the first year, then it will grow at the rate g and it will eventually be infinitely large. You get one pound the first year, you get 1+g pounds the second year, etc., $(1+g)^2$ the third year and so on. The present value of this--suppose it pays--let's say it pays c pounds each year, so it would be c times this. It would be c times $(1+g)^3$ in the third

year, etc., Then the present value is equal to $c/(r-g)$--that's the formula for the value of a growing console. g has to be less than r for this to make sense because if g--if it's growing faster than the rate of interest, then this infinite series will not converge and the value would be infinite. You might ask, "Well then how does that make sense?" What if the British Government promised to pay 10% more each year, how would the market value that? The formula doesn't have a number. I'll tell you why it doesn't have a number: the British Government will never promise to pay you 10% more each year because they can't do it. And, the market wouldn't believe them because you can't grow every year faster than the interest rate. Now that's one of the most basic lessons, you can't do it. One more thing that I think would be relevant to the--there's also the annuity formula. This is a formula that applies to--what if an asset pays a fixed amount every period and then stops? That's called an annuity. An annuity pays c dollars starting in $t = 1, 2, 3,$ and n is the last period, then it stops. A good example of an annuity is a mortgage on a house. When you buy a house, you borrow the money and you pay it back in fixed--it would usually be monthly, but let's say annual payments. You pay every year a fixed amount on your house to the mortgage originator and then after so many--n is 30 years, typically--you would then have paid it off. It used to be that mortgages had what's called a balloon payment at the end. This means that you would have to pay extra money at the end; but they decided that people have trouble doing that. It's much better to pay a fixed payment and then you're done. Otherwise, if you ask them to pay more at the end, then a lot of people won't have the money. We now have annuity mortgages. What is the present value of an annuity? That is, the present value of an annuity is equal to the amount--what did I say--$c*\{1 - [1/(1+r)]n\}/r$. That is the present value of an annuity.

I wanted to say one more thing because I realize that you have to--your first problem set will cover this--is to talk about the concept that applies probability theory to Economics. That is expected utility theory. Then I'll conclude with this. In Economics, it is assumed that people have a utility function, which represents how happy they are with an outcome--we typically take that as U, If I have a monetary outcome, then I have a certain amount of money, x dollars. How happy I am with x dollars is called $U(x)$. This, I think you've gotten from other economics courses--we have something called diminishing marginal utility. The idea is that for any amount of money--if this x is the amount of money that I receive--utility as the function of the amount of money I receive is downwardly-concave. The exact shape of the curve is subject to discussion, but the point of diminishing marginal utility is that, as you get more and more money, the increment in utility for each extra dollar diminishes. Usually we say it never goes down, we don't have it going down, cross that out. That would be where having more money makes you less happy. That may actually work that way, but our theory says no, you always want more. It's always upward sloping, but it may, after awhile, you get close to satiation where you've got enough.

Incidentally, I mentioned this last time--I was talking about--I was philosophizing about wealth and I asked what are you going to do with a billion dollars. We have many billionaires in this country and I think that the only thing you have to do with it is philanthropy. They have to give it away because they are essentially satiated. Because, as I said, you can only drive one car at a time and if you've got ten of them in the garage, then it doesn't really do you much good. You can't do it; you can't enjoy all ten of them. It's important--that's one reason why we want policies that encourage equality

of incomes--not necessarily equality, but reasonable equality--because the people with very low wealth have a very high marginal utility of income and people with very high wealth have very little. So, if you take from the rich and give to the poor you make people happier. We're not going to do that in a Robin Hood way; but in finance we're going to do that in a systematic way through risk management.

We're going to be taking away from lucky--you think of yourself as randomly on any point of this. You don't want--you know that you'd like to take money away from yourself in the high-outcome years and give it to yourself in the low-income years. What finance theory is based on--and much of economics is based on--the idea that people want to maximize the expected utility of their wealth. Since this is a concave function, it's not just the expected value. To calculate the expected utility of your wealth, you might also have to look at the expected return, or the geometric expected return, or the standard deviation. Or you might have to look at the fat tail. There are so many different aspects that we can get into and this underlying theory motivates a lot of what we do. But it's not a complete theory until we specify the utility function. Of course, we will also be talking about behavioral finance in this course and we'll, at times, be saying that the utility function concept isn't always right--the idea that people are actually maximizing expected utility might not be entirely accurate. But, in terms of the basic theory, that's the core concept.

I have one question on the problem set that asks you to think about how you would handle a decision: whether to gamble, based on efficient--based on expected utility theory. That's a little bit of a tricky question but--So, do the best you can on it and think--try to think about what this kind of theory would imply for gambling behavior. I will see you on Friday. That's two days from now in this room.

[end of transcript]